

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (previously presented) A method for maintaining a system for database management, the method comprising:
 - during splitting of a leaf block of a database index recording an address of a newly created leaf block;
 - maintaining the address of the newly created leaf block in a list as part of metadata of a primary B+tree; and
 - maintaining a measure of invalid guess-database block addresses by calculating a ratio of a count of database block addresses in the list of new addresses to a total number of leaf blocks of the primary B+tree.

2. (cancelled)

3. (previously presented) The method according to claim 1, wherein the measure of invalid guess-database block addresses applies to mapping tables and secondary indexes on the primary B+tree.

4. (previously presented) The method according to claim 1, wherein the list of database block addresses and the ratio are maintained only when the ratio is less than a threshold value.

5. (original) The method according to claim 4, wherein the threshold value for the ratio is about 10%.

6. (original) The method according to claim 3, further comprising:
adjusting a guess-DBA quality of at least one of the mapping table and the secondary index utilizing the ratio.

7. (original) The method according to claim 4, wherein if the ratio is below the threshold value the method further comprises:

selectively correcting entries in the mapping table and/or secondary index.

8. (original) The method according to claim 7, wherein correcting entries in the mapping table comprises for all rows in a list of blocks in the primary B+tree:
obtaining corresponding mapping table row identifiers and database block addresses of a current block in the list;

sorting the corresponding mapping table row identifiers;

obtaining mapping table rows corresponding to the mapping table row identifiers; and

updating a guess-DBA component if it has changed.

9. (original) The method according to claim 8, wherein the correcting is carried out on-line in a piece-wise manner.

10. (original) The method according to claim 4, wherein correcting entries in the secondary index comprises for all rows in a list of blocks in the primary B+tree:

obtaining a secondary index key, a primary key and a database block address of a current block in the list of blocks;

sorting the secondary index keys, primary keys and database addresses in order of (secondary index key, primary key) pairs;

obtaining an index row corresponding to the (secondary index key, primary key) pair; and

updating a guess-DBA component of the index row if the guess-DBA has changed.

11. (previously presented) The method according to claim 10, wherein the correcting is carried out on-line in a piece-wise manner.

12. (original) The method according to claim 4, wherein if the ratio is above the threshold value the method further comprises:

correcting guess-database addresses on a per object basis.

13. (original) The method according to claim 12, wherein correcting guess-database block addresses on the mapping table comprises:

performing a full scan of the mapping table;

determining for each row of the mapping table a correct guess-database block address by traversing the primary B+tree up to a penultimate level;

updating each row of the mapping table with the correct guess-database block address; and

committing the correct guess-database address to the mapping table in batches.

14. (original) The method according to claim 12, wherein correcting guess-database block addresses on a per object basis comprises for each secondary index object:

performing a full scan of the secondary index object;

determining for each row of the secondary index a correct guess-database block address by traversing the primary B+tree up to a penultimate level;

updating each row of the secondary index with the correct guess-database block address; and
committing the correct guess-database block address to the secondary index in batches.

15. (original) The method according to claim 1, further comprising:
maintaining a list of database block addresses in the list.

16. (previously presented) A system comprising:
a list of addresses of blocks newly created during splitting of a primary B+tree;
a count of database block addresses in the list; and
a ratio of count of database block addresses to total number of leaf blocks as a measure of invalid guess-database block addresses.

17. (cancelled)

18. (cancelled)

19. (original) The system according to claim 16, wherein the database index is a primary B+tree structure, wherein the system further comprises:

a mapping table used to support bitmap indexes.

20. (original) The system according to claim 19, further comprising:

a bitmap index supported by the mapping table.

21. (original) The system according to claim 16, wherein the database index is a primary B+tree structure, wherein the system further comprises:

a secondary index structure comprising hybrid row identifiers.

22. (previously presented) A computer program product for performing a process for maintaining a database management system, comprising:

a computer readable medium; and

computer program instructions, recorded on the computer readable medium, executable by a processor, for performing the steps of:

 during splitting of a leaf block of a primary B+tree recording an address of a newly created leaf block;

 maintaining the address of the newly created leaf block in a list as part of primary B+tree metadata; and

 maintaining a measure of invalid guess-database block addresses by calculating a ratio of a count of database block addresses in the list of new addresses to a total number of leaf blocks of the primary B+tree.

23. (previously presented) A system for performing a process for maintaining a database management system, comprising:

 a processor operable to execute computer program instructions; and

 a memory operable to store computer program instructions executable by the processor, for performing the steps of:

 during splitting of a leaf block of a primary B+tree recording an address of a newly created leaf block;

 maintaining the address of the newly created leaf block in a list as part of primary B+tree metadata; and

 maintaining a measure of invalid guess-database block addresses by calculating a ratio of a count of database block addresses in the list of new addresses to a total number of leaf blocks of the primary B+tree.

24. (new) The computer program product according to claim 22, wherein the measure of invalid guess-database block addresses applies to mapping tables and secondary indexes on the primary B+tree.

25. (new) The computer program product according to claim 22, wherein the list of database block addresses and the ratio are maintained only when the ratio is less than a threshold value.

26. (new) The computer program product according to claim 25, wherein the threshold value for the ratio is about 10%.

27. (new) The computer program product according to claim 24, further comprising:

adjusting a guess-DBA quality of at least one of the mapping table and the secondary index utilizing the ratio.

28. (new) The computer program product according to claim 25, wherein if the ratio is below the threshold value the computer program product further comprises:

selectively correcting entries in the mapping table and/or secondary index.

29. (new) The computer program product according to claim 28, wherein correcting entries in the mapping table comprises for all rows in a list of blocks in the primary B+tree:

obtaining corresponding mapping table row identifiers and database block addresses of a current block in the list;

sorting the corresponding mapping table row identifiers;

obtaining mapping table rows corresponding to the mapping table row identifiers; and

updating a guess-DBA component if it has changed.

30. (new) The computer program product according to claim 29, wherein the correcting is carried out on-line in a piece-wise manner.

31. (new) The computer program product according to claim 25, wherein correcting entries in the secondary index comprises for all rows in a list of blocks in the primary B+tree:

obtaining a secondary index key, a primary key and a database block address of a current block in the list of blocks;

sorting the secondary index keys, primary keys and database addresses in order of (secondary index key, primary key) pairs;

obtaining an index row corresponding to the (secondary index key, primary key) pair; and

updating a guess-DBA component of the index row if the guess-DBA has changed.

32. (new) The computer program product according to claim 31, wherein the correcting is carried out on-line in a piece-wise manner.

33. (new) The computer program product according to claim 25, wherein if the ratio is above the threshold value the computer program product further comprises:

correcting guess-database addresses on a per object basis.

34. (new) The computer program product according to claim 33, wherein correcting guess-database block addresses on the mapping table comprises: performing a full scan of the mapping table;

determining for each row of the mapping table a correct guess-database block address by traversing the primary B+tree up to a penultimate level;

updating each row of the mapping table with the correct guess-database block address; and

committing the correct guess-database address to the mapping table in batches.

35. (new) The computer program product according to claim 33, wherein correcting guess-database block addresses on a per object basis comprises for each secondary index object:

performing a full scan of the secondary index object;

determining for each row of the secondary index a correct guess-database block address by traversing the primary B+tree up to a penultimate level; updating each row of the secondary index with the correct guess-database block address; and committing the correct guess-database block address to the secondary index in batches.

36. (new) The computer program product according to claim 22, further comprising:

maintaining a list of database block addresses in the list.

37. (new) The system according to claim 23, wherein the measure of invalid guess-database block addresses applies to mapping tables and secondary indexes on the primary B+tree.

38. (new) The system according to claim 23, wherein the list of database block addresses and the ratio are maintained only when the ratio is less than a threshold value.

39. (new) The system according to claim 38, wherein the threshold value for the ratio is about 10%.

40. (new) The system according to claim 37, further comprising:

adjusting a guess-DBA quality of at least one of the mapping table and the secondary index utilizing the ratio.

41. (new) The system according to claim 38, wherein if the ratio is below the threshold value the system further comprises:

selectively correcting entries in the mapping table and/or secondary index.

42. (new) The system according to claim 41, wherein correcting entries in the mapping table comprises for all rows in a list of blocks in the primary B+tree:

obtaining corresponding mapping table row identifiers and database block addresses of a current block in the list;

sorting the corresponding mapping table row identifiers;

obtaining mapping table rows corresponding to the mapping table row identifiers; and

updating a guess-DBA component if it has changed.

43. (new) The system according to claim 42, wherein the correcting is carried out on-line in a piece-wise manner.

44. (new) The system according to claim 38, wherein correcting entries in the secondary index comprises for all rows in a list of blocks in the primary B+tree:

obtaining a secondary index key, a primary key and a database block address of a current block in the list of blocks;

sorting the secondary index keys, primary keys and database addresses in order of (secondary index key, primary key) pairs;

obtaining an index row corresponding to the (secondary index key, primary key) pair; and

updating a guess-DBA component of the index row if the guess-DBA has changed.

45. (new) The system according to claim 44, wherein the correcting is carried out on-line in a piece-wise manner.

46. (new) The system according to claim 38, wherein if the ratio is above the threshold value the system further comprises:

correcting guess-database addresses on a per object basis.

47. (new) The system according to claim 46, wherein correcting guess-database block addresses on the mapping table comprises:

performing a full scan of the mapping table;
determining for each row of the mapping table a correct guess-database block address by traversing the primary B+tree up to a penultimate level;
updating each row of the mapping table with the correct guess-database block address; and
committing the correct guess-database address to the mapping table in batches.

48. (new) The system according to claim 46, wherein correcting guess-database block addresses on a per object basis comprises for each secondary index object:

performing a full scan of the secondary index object;
determining for each row of the secondary index a correct guess-database block address by traversing the primary B+tree up to a penultimate level;
updating each row of the secondary index with the correct guess-database block address; and
committing the correct guess-database block address to the secondary index in batches.

49. (new) The system according to claim 23, further comprising:
maintaining a list of database block addresses in the list.

50. (new) A method of maintaining a system for database management, the method comprising:

modifying a primary index in the database management system;
storing information relating to the modification of the primary index in a secondary index or a mapping table in the database management system; and
maintaining a measure of quality of the secondary index or the mapping table based on a ratio of information relating to the modification of the primary index to a size of the primary index.

51. (new) The method of claim 50, wherein maintaining a measure of quality of the secondary index or a mapping table comprises calculating a ratio of the amount of information relating to the modification of the primary index in the secondary index or the mapping table to a size of the primary index.

52. (new) The method of claim 50, wherein modifying the primary index comprises splitting a leaf block of the primary index.

53. (new) The method of claim 52, wherein storing information relating to the modification of the primary index in a secondary index or a mapping table

comprises storing an address of a leaf block created during the splitting of the leaf block of the primary index in the secondary index or the mapping table.

54. (new) The method of claim 53, wherein maintaining a measure of quality of the secondary index or a mapping table comprises calculating a ratio of a number of addresses of leaf blocks created during the splitting of the leaf block of the primary index that are stored in the secondary index or the mapping table to a total number of leaf blocks in the primary index.

55. (new) The method of claim 50, further comprising:
adjusting a quality of at least one of the secondary index or the mapping table utilizing the ratio.

56. (new) The method of claim 50, wherein the ratio is maintained only when the ratio is less than a threshold value.

57. (new) The method of claim 56, wherein the threshold value for the ratio is about 10%.

58. (new) The method of claim 56, wherein if the ratio is below the threshold value the method further comprises:

selectively correcting the information in the secondary index or the mapping table.

59. (new) A computer program product for performing a process for maintaining a database management system, comprising:

 a computer readable medium; and
 computer program instructions, recorded on the computer readable medium, executable by a processor, for performing the steps of:
 modifying a primary index in the database management system;
 storing information relating to the modification of the primary index in a secondary index or a mapping table in the database management system; and
 maintaining a measure of quality of the secondary index or the mapping table based on a ratio of information relating to the modification of the primary index to a size of the primary index.

60. (new) The computer program product of claim 59, wherein maintaining a measure of quality of the secondary index or a mapping table comprises calculating a ratio of the amount of information relating to the modification of the primary index in the secondary index or the mapping table to a size of the primary index.

61. (new) The computer program product of claim 59, wherein modifying the primary index comprises splitting a leaf block of the primary index.

62. (new) The computer program product of claim 61, wherein storing information relating to the modification of the primary index in a secondary index or a mapping table comprises storing an address of a leaf block created during the splitting of the leaf block of the primary index in the secondary index or the mapping table.

63. (new) The computer program product of claim 62, wherein maintaining a measure of quality of the secondary index or a mapping table comprises calculating a ratio of a number of addresses of leaf blocks created during the splitting of the leaf block of the primary index that are stored in the secondary index or the mapping table to a total number of leaf blocks in the primary index.

64. (new) The computer program product of claim 59, further comprising: adjusting a quality of at least one of the secondary index or the mapping table utilizing the ratio.

65. (new) The computer program product of claim 59, wherein the ratio is maintained only when the ratio is less than a threshold value.

66. (new) The computer program product of claim 65, wherein the threshold value for the ratio is about 10%.

67. (new) The computer program product of claim 65, wherein if the ratio is below the threshold value the computer program product further comprises:
selectively correcting the information in the secondary index or the mapping table.

68. (new) A system for performing a process for maintaining a database management system, comprising:
a processor operable to execute computer program instructions; and
a memory operable to store computer program instructions executable by the processor, for performing the steps of:

modifying a primary index in the database management system;
storing information relating to the modification of the primary index in a secondary index or a mapping table in the database management system; and
maintaining a measure of quality of the secondary index or the mapping table based on a ratio of information relating to the modification of the primary index to a size of the primary index.

69. (new) The system of claim 68, wherein maintaining a measure of quality of the secondary index or a mapping table comprises calculating a ratio of the amount of information relating to the modification of the primary index in the secondary index or the mapping table to a size of the primary index.

70. (new) The system of claim 68, wherein modifying the primary index comprises splitting a leaf block of the primary index.

71. (new) The system of claim 70, wherein storing information relating to the modification of the primary index in a secondary index or a mapping table comprises storing an address of a leaf block created during the splitting of the leaf block of the primary index in the secondary index or the mapping table.

72. (new) The system of claim 71, wherein maintaining a measure of quality of the secondary index or a mapping table comprises calculating a ratio of a number of addresses of leaf blocks created during the splitting of the leaf block of the primary index that are stored in the secondary index or the mapping table to a total number of leaf blocks in the primary index.

73. (new) The system of claim 68, further comprising:

adjusting a quality of at least one of the secondary index or the mapping table utilizing the ratio.

74. (new) The system of claim 68, wherein the ratio is maintained only when the ratio is less than a threshold value.

75. (new) The system of claim 74, wherein the threshold value for the ratio is about 10%.

76. (new) The system of claim 74, wherein if the ratio is below the threshold value the system further comprises:

selectively correcting the information in the secondary index or the mapping table.